EMBEDDED ARTIFICIAL INTELLIGENCE ROBOT FOR FINANCIAL LITERACY: AN ANALYSIS OF THE SOCIAL IMPLICATIONS OF ARTIFICIAL INTELLIGENCE IN EDUCATION

A Research Paper submitted to the Department of Engineering and Society
Presented to the Faculty of the School of Engineering and Applied Science
University of Virginia • Charlottesville, Virginia
In Partial Fulfillment of the Requirements for the Degree
Bachelor of Science, School of Engineering

By

Sophie Howle

March 30, 2023

On my honor as a University student, I have neither given nor received unauthorized aid on this assignment as defined by the Honor Guidelines for Thesis-Related Assignments.

ADVISOR
Catherine D. Baritaud, Department of Engineering and Society

ARTIFICIAL INTELLIGENCE SHOWS PROMISE TO BE A VALUABLE EDUCATIONAL TOOL TO AID IN TEACHING IMPORTANT TOPICS.

Artificial intelligence (AI) is being used increasingly across the world, with many applications ranging from grocery store self-checkout systems to criminal defense algorithms. Nearly a $100 billion industry already, the global AI market is expected to make great strides, expanding at an estimated compound annual growth rate (CAGR) of over 38% by 2030 (Grand View Research, n.d.). Due to its wide number of applications and versatility, AI can and is being used across virtually any discipline. According to research conducted by Su & Yang (2022), AI tools are increasingly being used in the early childhood education (ECE) sphere for a variety of different reasons, a primary motivation of which is to help improve learning by increasing engagement and offering personalized learning experiences.

The aid provided to educators from AI systems offers the potential to act as a welcomed lifeline in the sea of hardships they currently face. Education in the U.S. is far from perfect, with American students' achievement lagging significantly behind that of their counterparts in many other developed nations (DeSilver, 2019). There exist issues within U.S. education already that can have lasting effects on students. For example, racial and gender bias towards children in modern education is pervasive and well documented (Anokam et al., 2022). Even within important, skill-building topics, educators struggle to maintain equity and avoid bias. According to 2018 research by Marri & Sonu, there exists specific "hidden curricula" within existing financial literacy lessons. Despite the importance of these lessons for future economic success, they tend to require understanding and/or acceptance of neoclassical economics, which largely fails to acknowledge the complex history of economic inequality (Marri & Sonu, 2018).

Additionally, grade schools are in the midst of a burnout crisis, with K-12 educators reporting the highest burnout rate of all U.S. professions (Agrawal & Marken, 2022). During the height of the COVID-19 pandemic, educators were more likely to report increased levels of stress, anxiety, and burnout, and the effects did not stop with the end of social distancing protocols. There have been lasting effects of staff shortages and a resulting increased workload (Walker, 2021). The burnout crisis does not only affect teachers, but also can have implications for students. Without effective, well-equipped teachers, learning can be impacted.

There is hope that AI systems can address some of the hardships of childhood education, primarily by reducing the burden on teachers. Moreover, some experts believe children should be engaging with an AI curriculum as a key aspect of modern digital literacy lessons in a rapidly developing digital world (Yang, 2022). Research done by Druga et al. (2019) shows that the already-existing digital divide between racial minority and dominant groups has been deepened in the realm of AI technologies. These researchers claim there is an urgent need for the promotion of AI in the classroom to attain digital equity and close the observed divide.

To further explore the intricacies of AIEd, this paper will seek to answer the question of whether artificially intelligent systems should be implemented in K-12 education programs by exploring the social consequences and discussing a case study on the implementation of one AIEd prototype. This will be done by first contextualizing the problem at hand: that AIEd systems are being implemented without adequate insight into their effects on students. This issue serves as the primary motivation behind my STS research and will be explored through the lens of this research. An overview of the two, tightly coupled STS and technical research projects will then be provided in the context of the aforementioned research question. Details on the objective, metrics and methodology of the technical project will then be provided, as well as an analysis of

the system under the Technology and Social Relationships STS framework. A brief summary of the results of this project will then be presented in order to synthesize this case study with the related STS research, allowing me to make final conclusions and resulting recommendations about the future of AIEd systems. Both the technical project and STS research were carried out over the course of the 2022-2023 academic year, totaling 8 months' time. The technical project was constrained by the funding and objectives provided by our partners at The MITRE Corporation. The research for this project involved experimental testing with a sample size of 26 participants, all of which were fourth grade students at the same elementary school in Charlottesville, VA. While the time and budget constraints led us to select this relatively limited group of participants, additional research across varying age groups, geographic locations, language abilities, and other factors is ultimately recommended.

ARTIFICIAL INTELLIGENCE SYSTEMS ARE BEING IMPLEMENTED IN EDUCATION PROGRAMS DESPITE A LACK OF ADEQUATE INSIGHT INTO THE POTENTIAL CONSEQUENCES.

In response to the growing popularity of AI, as well as the promise of these systems to bring about beneficial changes, AI systems are starting to be implemented within all spheres of education, including grade school. AI in education (AIEd) shows promise to address issues of teacher burnout and improve learning. However, there are relatively minimal amounts of research regarding AIEd compared to the level at which it has been implemented. A 2021 study by researchers Aslan and Zhang found only 40 comprehensive, empirical journal articles written prior to 2020 regarding AIEd from a search of three major databases. Despite this stark lack of

research, 92% of academic institutions have implemented such systems (University of the People, n.d.).

Moreover, there is not only an insufficiency of evidence proving the benefits of AI systems for ECE, but there is also a potentially more concerning insufficiency of evidence exploring the potential harms of such systems (Aslan & Zhang, 2021). Many advocates against the widespread use of AI warn of the potential for embedding subconscious bias into these systems, and therefore reinforcing and amplifying existing structures of inequity (Kantayya, 2020). Unfortunately, evidence of these embedded biases and the real-world consequences has already begun to surface. A 2016 study by Angwin et al. found that a criminal justice algorithm employed in Florida two years prior disproportionately mislabeled Black defendants as "high risk" at almost twice the rate of their White counterparts, making them more likely to receive harsher sentences and serve more jail time. A likely culprit contributing to these biases is the relative lack of diversity amongst AI programmers. For example, less than 3% of Google's entire workforce and 4% of Facebook's and Microsoft's are Black. The story is similar in terms of gender, with over 75% of AI professionals identifying as male (Hao, 2019).

AIEd is not immune to the same fate as the many already observed instances of biases being cemented into algorithms and, therefore, worsening existing disparities against marginalized groups. Despite promises and hopes alike that AI can level the playing field in the classroom between minority students and their classmates of dominant groups, the potential for these systems to exacerbate the already existing divide proves to be a larger threat. If education becomes too digitized without properly considering how to check the inherent biases of the predominantly White developers of these systems, the unforeseen consequences can have disastrous and lasting effects.

SHOULD ARTIFICIALLY INTELLIGENT SYSTEMS BE IMPLEMENTED WITHIN K-12 EDUCATION PROGRAMS?

To address the growing demand for such systems and corresponding need for related research, my senior project, involving a technical capstone project and technical paper and a related Science, Technology, and Society (STS) research paper, addresses the ethics, practicality, and implementation of artificially intelligent systems used in education. The former project involves the logistics and practice of implementation, including building an AI system for financial literacy educational usage, embedding the system into a humanoid robot, and testing with human subjects. The latter project involves researching and reporting on the societal implications of such projects, especially regarding potential bias in AI systems and the competency of the existing research on the subject matter, as seen in the previous and following paragraphs. The two projects are tightly coupled, as the latter investigates the social consequences of the former in an effort to understand whether implementing these systems on a large scale is an appropriate course of action.

The technical project addresses the question of whether the embedded AI system we created provides a benefit to students. The objective of this project is to create a system adequately equipped to teach financial literacy lessons to grade school students, to test the system and make appropriate adjustments, and to use the experimental findings to ultimately determine if the system's implementation provides a benefit to students. The metrics across which the project's hypothesis will be tested are as follows in Figure 1: Technical Project Metrics Tree.

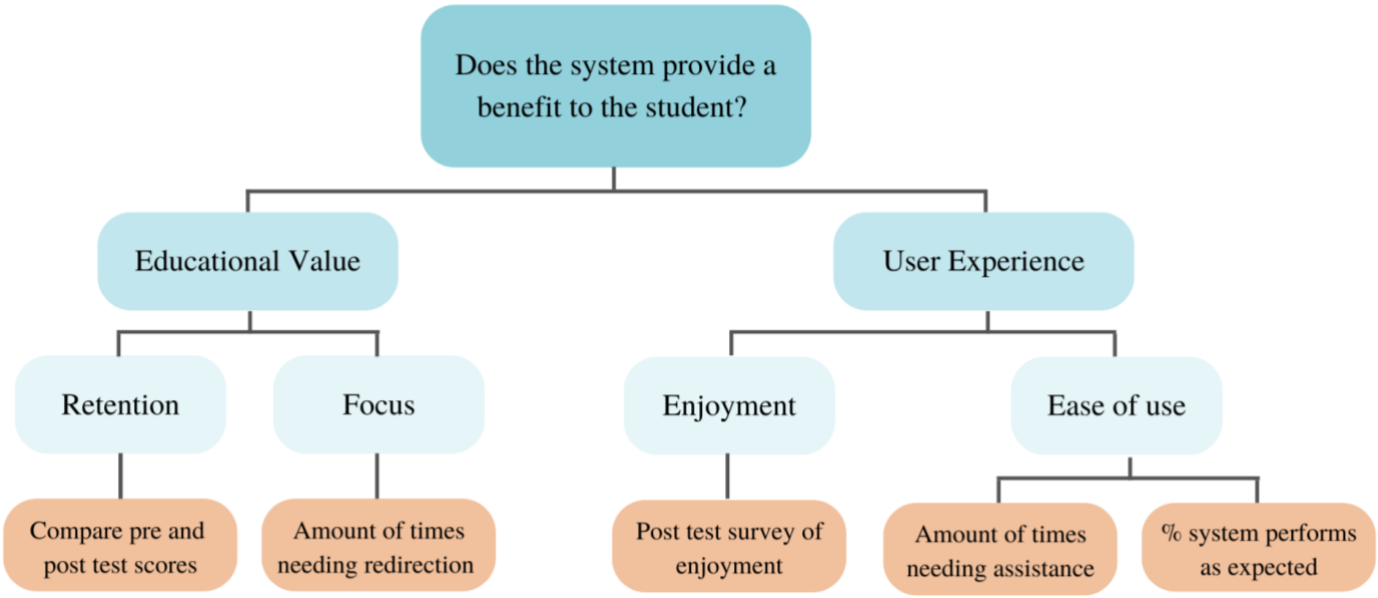

Figure 1: Technical Project Metrics Tree: This diagram shows the breakdown of the objective of our technical project, the associated metrics, and the data that will be collected to measure these metrics (Howle, 2022).

The methods used to answer the question at hand involved both implementation and testing. The implementation stage involves building the AI system equipped with financial literacy lessons and embedded this system into a NAO robot, a small, programmable humanoid robot designed for interacting with humans, making it ideal for usage within the classroom (Institute of Electrical and Electronics Engineers, n.d.). Testing involves comparing the performance of students who learned financial literacy lessons via our robot with the performance of those who did not have access to the robot. Those who do not have access to the robot will learn the same lesson from a guided worksheet. After analyzing the difference in scores between the two groups, as well as data regarding the performance of the system and the enjoyment of the students, the research team will make a conclusion and corresponding recommendation regarding whether the system is beneficial and should be further implemented

and invested in. An aspect of the aforementioned technical project that was not considered was any training or research into how to ensure our system does not hold embedded racial or gender biases. I found the lack of interdisciplinary perspectives concerning, especially for a system designed to be used with as sensitive a group as young children.

In response, this STS research paper seeks to answer the question of what the potential social consequences of implementing artificial intelligence in K-12 education are. The research will explore how both students and educators stand to be impacted by the implementation of AI systems in education, such as that of my technical project. An analysis of the social implications in terms of the major stakeholders and the ways in which they stand to be impacted by implementation of these systems - using the Technology and Social Relationships STS framework - is demonstrated in Figure 2: Technology and Social Relationships Model.

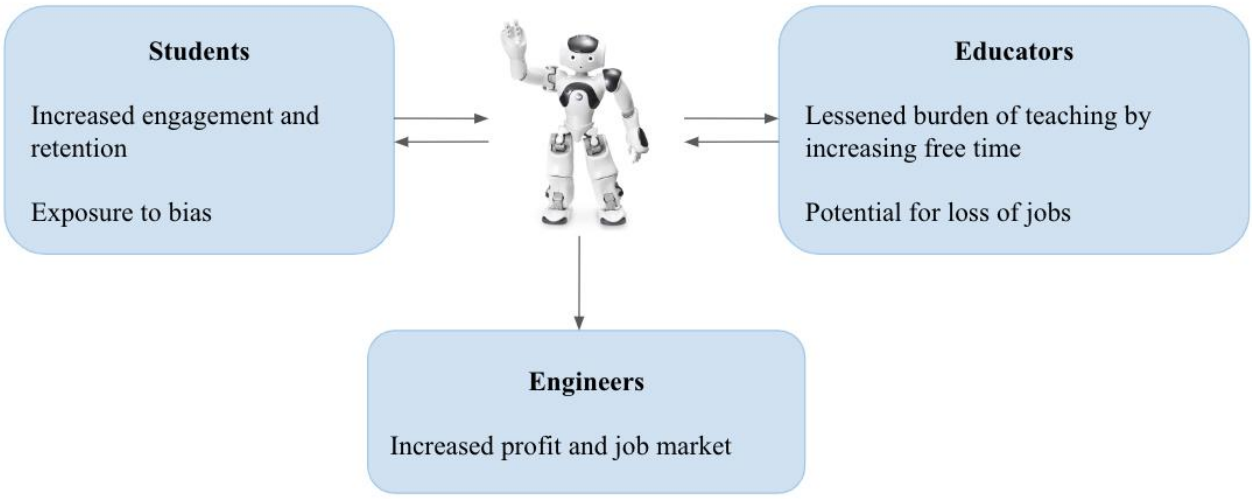

Figure 2: Technology and Social Relationships Model: An analysis of the potential benefits and harms to be experienced by stakeholders in the implementation of AI systems in education. In this figure, an arrow pointing towards a stakeholder indicates a potential benefit to that group, while an arrow pointing away from a stakeholder indicates a potential harm to that group (Howle, 2023).

Arguably the most important factor to consider in analyzing AIEd systems is the effect on students. The aforementioned research on how AIEd holds potential to benefit students by improving their learning has created demand for such systems and has led to their implementation already (Su & Yang, 2022). However, with racial and gender bias towards children in both modern education and existing algorithms being pervasive and well documented, and without proper input from education experts on these systems, the potential for bias to be embedded into these systems and imposed onto young students stands as a major concern (Anokam et al., 2022; Su & Yang, 2022). Due to the sensitive nature of the development and learning of children, some experts argue that the potential for harm outweighs that of any potential benefits (Perry & Lee, 2019).

Another important stakeholder who stands to be affected by the systems in question is educators. AI poses a potential benefit of lessening the burden on teachers in a number of ways, which could aid in solving the current burnout crisis. The ways in which teachers stand to save time and energy as a result of AI being implemented in the classroom include automated grading, feedback loops for student evaluations, and providing virtual facilitators for simpler tasks and lessons (Owoc et al., 2021). However, while teachers stand to benefit from AI, there is also fear surrounding the potential loss of jobs. In fact, British education expert and Vice-Chancellor of the University of Buckingham Anthony Seldon predicts that robots will replace teachers by 2027 (The University of Buckingham, 2017).

The final major stakeholder to consider is the engineers developing these AIEd systems. Those developing these systems almost exclusively stand to benefit from their implementation, as it results in both increased profit and job security. Without facing a major threat of harm, developers of AI systems are likely to promote and encourage their deployment, oftentimes even

without proper checks in place (Kantayya, 2020). Especially considering the harm done to the students on which they are implemented may not become obvious until years in the future when longitudinal studies are completed, the interests of the engineers as a stakeholder should not be prioritized above that of students and educators.

REGARDLESS OF WHETHER MY TECHNICAL PROJECT SHOWS A BENEFIT TO STUDENTS, ANY POTENTIAL HARM SHOULD BE TAKEN INTO SERIOUS CONSIDERATION BEFORE SYSTEMS ARE IMPLEMENTED.

Testing my technical project in elementary school students thus far has shown a need for improvements to the system, and therefore, the system is deemed inadequate for widespread implementation in its current state. Once improvements have been made, the system can be re-tested and re-evaluated for its benefit to students. However, as a result of the findings of this STS research paper, I recommend waiting to implement this system even if it shows a benefit to students with further development and experimentation, as the experiments do not test for potential harm - only for educational benefits. Additionally, I recommend introducing more interdisciplinary perspectives in the work, including experts in both childhood education and bias prevention. Another strategy to improve upon the system would be to ensure the teams employed to work on this project are diverse in racial and gender identity.

This corresponding research and analysis in my STS paper reveals a need for adjustments to current and future systems, as the existing research on their implications is inadequate compared to the potential consequences. Although current research mostly shows only the *potential* for harm to students, rather than guaranteed harm, the recommendations should reflect the worst case scenario in order to prevent any irreversible harm to marginalized groups. Further,

the analysis above reveals a need for more research on the impacts of AIEd; most importantly, in identifying the potential for harm and the likelihood of such occurrences. Within this research, strategies should be identified for overcoming inherent biases to ensure the AIEd systems we employ are equitable and overwhelmingly beneficial to both students and educators.

In conclusion, the interests of sensitive and marginalized groups must be held in highest regard, taking special consideration into the potential for harm to these groups and employing all strategies to avoid such consequences. AI systems are exciting, new technologies sure to disrupt society as we know it. However, with uncharted territory often comes a large margin for error. In this case, because the largely unknown and potentially dangerous consequences will impact children, it is more important than ever to proceed with great caution to keep our society's young students out of harm's way and make strides towards a more equitable society.

REFERENCES

Agrawal, S. & Marken, S. (2022, June 13). K-12 Workers Have Highest Burnout Rate in U.S. *Gallup.* https://news.gallup.com/poll/393500/workers-highest-burnout-rate.aspx

Anderson, M., Funk, C., Rainie, L., & Tyson, A. (2022, March 17) *AI and Human Enhancement: Americans' Openness is Tempered by a Range of Concerns.* Pew Research Center.

Angwin, J., Kirchner, L., Larson, J., & Mattu, S. (2016, May 23). Machine Bias: There's software used across the country to predict future criminals. And it's biased against blacks. *ProPublica.* https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing

Anokam, C., Blackson, E. A., Gerdes, M., Johnson, T. J., & Segan, E. (2022, April 14). Racial bias toward children in the early childhood education setting. *Journal of Early Childhood Research*, 20(3), 277-292. https://doi.org/10.1177/1476718X22108705

Aslan, A. B. & Zhang, K. (2021). AI technologies for education: Recent research & future directions. *Computers and Education: Artificial Intelligence, 2*(1). https://doi.org/10.1016/j.caeai.2021.100025.

Blackwell, C. K., Lauricella, A. R., & Wartella, E. (2014, August). Factors influencing digital technology use in early childhood education. *Computers & Education, 77*(1), 89-90. https://doi.org/10.1016/j.compedu.2014.04.013

Carlson, W. B. (2009). STS Frameworks: Technology Social Relationship Model. In *STS 4500: STS and Engineering Practice: Fall 2022* [Lecture notes]. School of Engineering and Applied Science, University of Virginia. Charlottesville, VA.

Council for Economic Education. (2021). *National Standards for Personal Finance Education.* https://www.councilforeconed.org/wp-content/uploads/2021/10/2021-National-Standards-for-Personal-Financial-Education.pdf

DeSilver, D. (2017, February 15). U.S. students' academic achievement still lags that of their peers in many other countries. *Pew Research Center.* https://www.pewresearch.org/fact-tank/2017/02/15/u-s-students-internationally-math-science/

Druga, S., Likhith, E., Qiu, T., Vu, S. T. (2019, March). Inclusive AI literacy for kids around the world. https://doi.org/10.1145/3311890.3311904

Gartner. (n.d.). Virtual Assistant. In *Gartner Information Technology Glossary*. Retrieved December 13, 2022 from https://www.gartner.com/en/information-technology/glossary/virtual-assistant-va

Gartner (2021, November 22). Gartner forecasts worldwide artificial intelligence software market to reach $62 billion in 2022 [Press release]. Retrieved from https://www.gartner.com/en/newsroom/press-releases/2021-11-22-gartner-forecasts-worldwide-artificial-intelligence-software-market-to-reach-62-billion-in-2022

Grand View Research. (n.d.). *Artificial Intelligence Market Size Report, 2022-2030: Report*

Overview. Grand View Research. Retrieved November 29, 2022 from https://www.grandviewresearch.com/industry-analysis/artificial-intelligence-ai-market

Guilherme, A. (2019). AI and education: the importance of teacher and student relations. *AI & Society*, 34(1), 47–54. https://doi.org/10.1007/s00146-017-0693-8

Hao, K. (2019, April 17). AI's white guy problem isn't going away. *MIT Technology Review.* https://www.technologyreview.com/2019/04/17/136072/ais-white-guy-problem-isnt-going-away/

Howle, S. (2022). *Technical Project Metrics Tree.* [Figure 1]. *STS Research Outline* (Unpublished undergraduate thesis). School of Engineering and Applied Science, University of Virginia. Charlottesville, VA.

Howle, S. (2023). *Technology and Social Relationships Model.* [Figure 2]. *STS Research Outline* (Unpublished undergraduate thesis). School of Engineering and Applied Science, University of Virginia. Charlottesville, VA.

Institute of Electrical and Electronics Engineers. (n.d.). *Nao*. Robots. https://robots.ieee.org/robots/nao/

Johnson, D.G. (n.d.) Social Construction of Technology. In *Encyclopedia of Science, Technology, and Ethics.* Macmillan Reference USA.

Kantayya, S. (Director). (2020). *Coded Bias* [Film]. 7th Empire Media.

Lusardi, A. (2019, January 24). Financial literacy and the need for financial education: evidence and implications. *Swiss Journal of Economics and Statistics*, 155(1), 1. https://doi.org/10.1186/s41937-019-0027-5

Marri, A. R. & Sonu, D. (2018). The hidden curriculum in financial literacy: Economics, standards, and the teaching of young children. *CUNY Academic Works: Publications and Research.* Retrieved December 1, 2022 from https://scholar.google.com/scholar?hl=en&as_sdt=0%2C33&q=financial+literacy+childhood+education+ethics&btnG=.

Owoc, M. L., Sawicka, A., & Weichbroth, P. (2021, August 6). Artificial intelligence technologies in education: Benefits, challenges and strategies of implementation. *Artificial Intelligence for Knowledge Management.* https://doi.org/10.1007/978-3-030-85001-2_4

Perry, A. M. & Lee, N. T. (2019, September 26). AI is coming to schools, and if we're not careful, so will its biases. *Brookings.* https://www.brookings.edu/blog/the-avenue/2019/09/26/ai-is-coming-to-schools-and-if-were-not-careful-so-will-its-biases/

Robotical. (n.d.) *The Benefits of Using Humanoid Robots in Education*. Robotical Ltd. Retrieved October 16, 2022, from https://robotical.io/wp-content/uploads/2021/08/Robotical_Benefits-of-Using-Humanoid-Robots.pdf.

School Robots by RobotLab. (n.d.) Retrieved from https://www.robotlab.com/classroom-robots

Su, J. & Yang, W. (2022). Artificial intelligence in early childhood education: A scoping review. *Computers and Education: Artificial Intelligence, 3*(1). https://doi.org/10.1016/j.caeai.2022.100049

The University of Buckingham. (2017, September 12). Robots will replace teachers in the next ten years, Vice-Chancellor reveals. *The University of Buckingham.* https://www.buckingham.ac.uk/news/robots-will-replace-teachers-in-the-next-ten-years-vice-chancellor-reveals/

University of the People. (n.d.) AI In Education: Where Is It Now And What Is The Future. *University of the People.* Retrieved December 1, 2022 from https://www.uopeople.edu/blog/ai-in-education-where-is-it-now-and-what-is-the-future/

Walker, T. (2021, November 12). Getting Serious About Teacher Burnout. *National Education Association News.* https://www.nea.org/advocating-for-change/new-from-nea/getting-serious-about-teacher-burnout

Yang, W. Artificial Intelligence education for young children: Why, what, and how in curriculum design and implementation. *Computers and Education: Artificial Intelligence, 3*(1). https://doi.org/10.1016/j.caeai.2022.100061